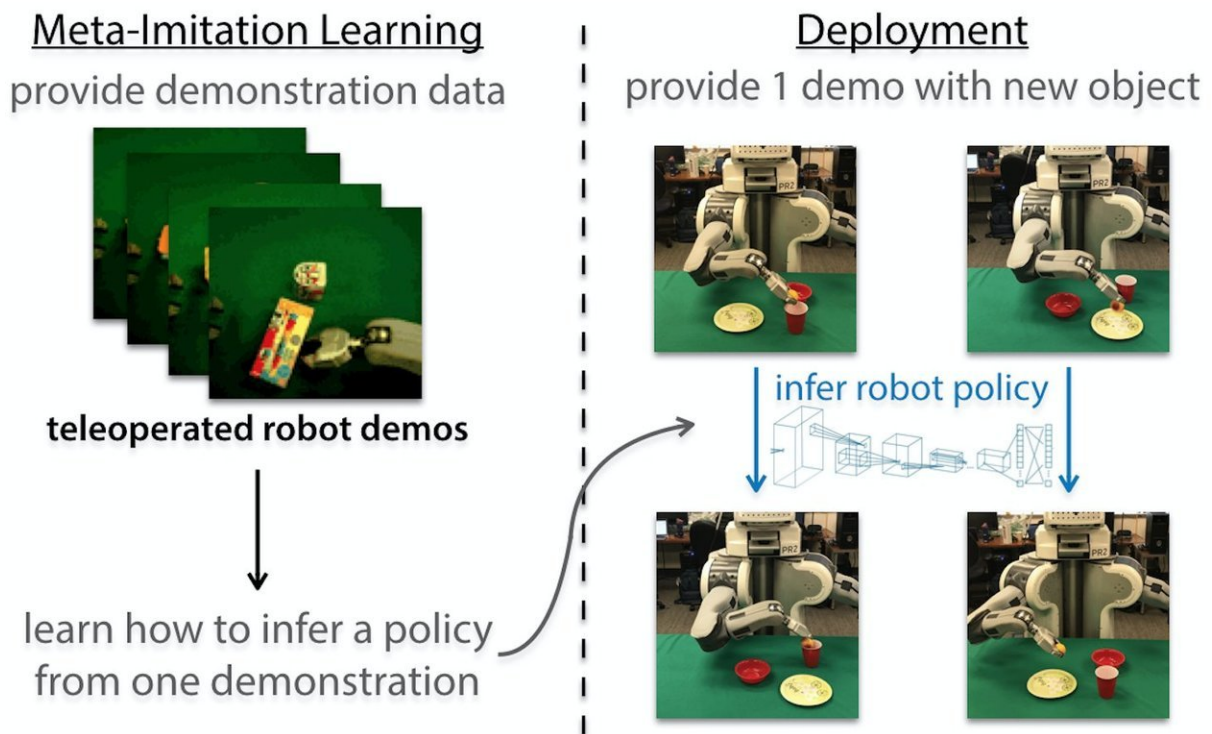


Robot able to mimic an activity after observing it just one time

July 4 2018, by Bob Yirka



Three steps for our meta-learning algorithm. Credit: Tianhe Yu and Chelsea Finn

A team of researchers at UC Berkeley has found a way to get a robot to mimic an activity it sees on a video screen just a single time. In a paper they have uploaded to the *arXiv* preprint server, the team describes the approach they used and how it works.

Robots that learn to do things simply by watching a human carry out an action a single time would be capable of learning many more new actions much more quickly than is now possible. Scientists have been working hard to figure out how to make it happen.

Historically though, robots have been programmed to perform actions like picking up an object by via code that expressly lays out what needs to be done and how. That is how most robots that do things like assemble cars in a factory work. Such robots must still undergo a training process by which they are led through procedures multiple times until they are able to do them without making mistakes. More recently, robots have been programmed to learn purely through observation—much like humans and other animals do. But such imitative learning typically requires thousands of observations. In this new effort, the researchers describe a technique they have developed that allows a [robot](#) to perform a desired action by watching a human being do it just a single time.

To accomplish this feat, the researchers combined imitation learning with a meta-learning algorithm. The result is something they call model-agnostic meta-learning (MAML). Meta-learning, the researchers explain, is a process by which a robot learns by incorporating prior experience. If a robot is shown video of a human picking up a pear or another similar object, for example, and putting it into a cup, bowl or other [object](#), it can get a "feel" for an objective. If in each instance it is taught to imitate the [behavior](#) in a certain way, then it "learns" what to do when observing other similar behaviors. Thus, when it sees a [video](#) of a person picking up a plum and putting it into a bowl, it recognizes the behavior and is able to translate that into a similar behavior of its own, which it can then carry out.

More information: — One-Shot Imitation from Observing Humans via Domain-Adaptive Meta-Learning, arXiv:1802.01557 [cs.LG] , arxiv.org/abs/1802.01557

Abstract

Humans and animals are capable of learning a new behavior by observing others perform the skill just once. We consider the problem of allowing a robot to do the same—learning from a raw video pixels of a human, even when there is substantial domain shift in the perspective, environment, and embodiment between the robot and the observed human. Prior approaches to this problem have hand-specified how human and robot actions correspond and often relied on explicit human pose detection systems. In this work, we present an approach for one-shot learning from a video of a human by using human and robot demonstration data from a variety of previous tasks to build up prior knowledge through meta-learning. Then, combining this prior knowledge and only a single video demonstration from a human, the robot can perform the task that the human demonstrated. We show experiments on both a PR2 arm and a Sawyer arm, demonstrating that after meta-learning, the robot can learn to place, push, and pick-and-place new objects using just one video of a human performing the manipulation.

— bair.berkeley.edu/blog/2018/06/28/dam/

© 2018 Tech Xplore

Citation: Robot able to mimic an activity after observing it just one time (2018, July 4) retrieved 8 April 2024 from <https://techxplore.com/news/2018-07-robot-mimic.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.
